


Technical Progress Report 11/1/94 - 1/31/95
Construction of a Connectionist Network Supercomputer
University of California, Berkeley
ONR URI Grant No. N00014-92-J-1617



1 Abstract


This report presents a summary of our proposal for years four and five of the project and the technical status for the period 11/1/94-1/31/95.

2 Plans for Years Fours and Five

This quarter we submitted a project renewal proposal. The proposal outlined our work for the first three years and presented the work remaining in the project. We summarize the proposal below.

2.1 Project Highlights

- A new digital vector microprocessor architecture has been invented and will soon be available in working systems. This architecture, called Torrent, is ideally suited to a variety of tasks requiring very high performance fixed point calculations.
- An infrastructure of software tools and techniques has been developed to support the new Torrent architecture. These tools are currently being used to test the chip design and also port applications code to the architecture.
- Much of the research involved with the VLSI project, including CAD tool improvements, circuits design and library development, has been incorporated in the university curriculum.
- Network interface test chips have been built to try new ideas for interconnecting processing nodes. A link bandwidth of 300 Mbits per second per wire has been achieved.
- Several generations of analog auditory preprocessor chips (smart sensors) are working in the lab. Recent design improvements will make it easier to assemble the chips into complete front end systems.
- A prototype speech recognition system using the Torrent simulator has been demonstrated. This represents a milestone in system design and integration.



19951031 026

- A locally developed object-oriented language, Sather, is finding wide acceptance in the research community. This work is being extended as pSather for use on parallel computing platforms, including CNS.
- Research Publications: 15 journal articles, 7 book chapters, 35 conference proceedings and 21 technical reports.
- Degrees: 3 Ph.D. theses and 3 Master's reports to date.

2.2 New work for years 4 and 5

In the final two years of the URI grant, we will focus on building systems based on the technology from the first three years, and on new science that is enabled by those systems. For instance, theoretical work on auditory or transition-based models for speech has led to algorithms that will require massive amounts of digital computation. Systems that we will be building in years three and four will be critical for basic research to be undertaken in years four and five. In another project the single-chip, low-power analog implementation of auditory spectral shape algorithm will be used as a building block for the construction of larger systems. These systems will enable in-depth speech recognition experiments and may lead to robust low-power single-chip speech recognizers.

This new work will then fulfill several goals that we have had from the start of the program:

1. To test our technology for high-performance computation on actual research problems, so that we can assess the utility of what we have done.
2. To accomplish new science that would not have been possible without our technological efforts.

We can divide the remaining work for years four and five into *Systems Development* and *Applications*. We will work on two different hardware platforms: 1) Vector Processor Systems, and 2) Analog VLSI Systems. Our application area for these systems as before will be speech recognition. We will continue a modest effort in vision processing.

The digital systems building activity for years four and five will be dominated by software work. However, there will be some hardware work required to develop systems with more than one node. There is still research and development work required, for instance, for the design and evaluation of an efficient network interface and router for our vector microprocessor. We also plan on one revision of the current chip that will make it more useful for multi-mode systems and for slower, larger memories. While we are currently planning on constraining these hardware projects to the minimum necessary to build multi-node systems, realistically it should be expected that some significant resources will be expended in this direction.

3 Technical Status

3.1 Software and Applications

High-level software. Progress continues to be excellent in the high-level software area of the project. The public release of Sather 1.0 has been successfully ported to a number of platforms and acceptance is very good. There are now hundreds of sites running the system and many are contributing to its development. A tutorial on Sather has been published both on-line and as a technical report. A preliminary version of a Sather interpreter was built and is being tested.

The parallel version, pSather, has also made excellent progress; a complete syntax checker for the language has been incorporated into the Sather 1.0 system and is in use. The run time system has been implemented and a proof of its safety properties completed. The group was invited to present a paper on pSather to the POOMA conference on parallel object-oriented languages and this was well received. There is now a complete system for a subset of the pSather 1.0 system and this will be extended to the complete design by summer of 1995. One of the major applications targets for Sather and pSather is the connectionist simulator ICSIM. Ben Gomes has completed one version of ICSIM and has started work on the parallel version. He will defend his thesis proposal in a qualifying exam early in 1995. This work makes heavy use of our previous studies of connectionist algorithms for parallel machines.

Low-level software. The main of emphasis of low level software development has been on further broadening of the SPERT libraries. Specifically:

- The core set of matrix operations for back propagation training have been incorporated into the library. In addition, non-vectorized (but bitwise-identical) portable versions of these routines have been written to aid applications development.
- More vectorized function approximations routines have been added to the library.
- A consistent set of basic vector integer operations has been built up.
- A set of fixed point FFT routines is being developed and debugged. It is intended that these routines will be used in our speech recognition front end system, RASTA.

In addition, a start was made on production testing and diagnostic software for the SPERT board.

Speech application. In the previous period, we demonstrated the use of fixed point Torrent libraries to implement a simplified neural network as part of a speech recognition application (BeRP). In the current period, we have been working on Torrent libraries to

implement a number of signal processing functions that are required for feature extraction on the speech. Towards this end, we have developed a number of routines for the computation of Fourier transforms, and have begun looking at the other pieces of the Perceptual Linear Prediction (PLP) algorithm in order to see what needs to be vectorized.

We also have continued our work on the SPAM and REMAP approaches to speech processing. These approaches will make recognition more robust and accurate but will require considerably more computing than we currently have available. We have begun a number of experiments with transition-based recognition of digits, and our initial results seem to confirm that we can make a simple system that will perform at least as well in noise as our more conventional systems while using roughly one-eighth of the frames for training. As we progress to more developed methods and increase training data, we expect the robustness to improve over the conventional phone-based approach. We also have developed the modifications to our object-oriented neural network training software that will be required to conduct tests of the recursive REMAP algorithm. Implementing these algorithms on large vocabulary continuous speech recognition will require the new hardware, but we can confirm many of the basic ideas now using the smaller digits task.

3.2 Hardware Development

Vector Processors. The SPERT board logic design was completed during last quarter, including a design review by an outside consultant. The circuit board has been laid out, and is now out for fabrication. Completed boards are expected back in early March.

Analog VLSI pre-processors. As described in the last status report we recently designed an extension to AER which permits many chips to send out neural representations on a shared bus efficiently, fabricated a small test chip using the extended protocol, and built a 7-sender demonstration system using the test chip.

Several advances have occurred over the past three months with regards to the extended AER protocol. A 128-channel spectral-shape auditory-preprocessor was fabricated using this extended AER protocol. These chips were tested and found to be functional, and a 3-chip system was constructed using three copies of this chip. Encouraged by this result, 24 additional copies of this chip are currently in fabrication, and we plan to build a system with a larger number of chips upon its receipt.

Our submission to the Advanced Research in VLSI conference was accepted (this competitive conference has a acceptance ratio of 25% percent). We finished an updated version of the 12-page paper, which will be published in a hard-bound book format and distributed at the conference.

Finally, our collaboration with Richard Lippmann at MIT Lincoln Labs on analog approaches to word-spotting is proceeding well, and we expect to report on initial results from these efforts in the next status report.

4 Recent Publications

[lazzaro1] Lazzaro, J. and Wawrzynek, J., "A Multi-Sender Asynchronous Extension to the AER protocol," accepted, 1995 Advanced Research in VLSI conference.

[Philippsen] Philippsen, M., "Sather 1.0 Tutorial," ICSI Technical Report TR-94-062, December 1994.

[Bregler] Bregler, C., and Omohundro, S., "Nonlinear Image Interpolation using Surface Learning," Neural Information Processing Systems Conf., Denver CO, Nov. 29-Dec. 1, 1994.

[feldman] Feldman, Lakoff, Bailey, Narayanan, Regier, and Stolcke, "L0-The First Five Years," to be published in *AI Review*, Vol. 8, special issue on Integration of Natural Language and Vision Processing, edited by Paul McKeivitt.

[Stoutamire] Feldman, J. and Stoutamire, D., "Language Support for Relaxed Consistency in pSather 1.0," submitted to conference on PPOP (Principles and Practice of Parallel Programming), Summer 1995.

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>per attached</i>	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
<i>A-1</i>	



OFFICE OF THE UNDER SECRETARY OF DEFENSE (ACQUISITION)
DEFENSE TECHNICAL INFORMATION CENTER
CAMERON STATION
ALEXANDRIA, VIRGINIA 22304-6145

IN REPLY
REFER TO

DTIC-OCC

SUBJECT: Distribution Statements on Technical Documents

TO:

OFFICE OF NAVAL RESEARCH
CORPORATE PROGRAMS DIVISION
ONR 353
800 NORTH QUINCY STREET
ARLINGTON, VA 22217-5660

- 1995 1031 024
1. Reference: DoD Directive 5230.24, Distribution Statements on Technical Documents, 18 Mar 87.
 2. The Defense Technical Information Center received the enclosed report (referenced below) which is not marked in accordance with the above reference.
PROGRESS REPORT
N00014-92-J-1617
TITLE: CONSTRUCTION OF A
CONNECTIONIST NETWORK
SUPERCOMPUTER
 3. We request the appropriate distribution statement be assigned and the report returned to DTIC within 5 working days.
 4. Approved distribution statements are listed on the reverse of this letter. If you have any questions regarding these statements, call DTIC's Cataloging Branch, (703) 274-6837.

FOR THE ADMINISTRATOR:

1 Encl

GOPALAKRISHNAN NAIR
Chief, Cataloging Branch

FL-171
Jul 93

DISTRIBUTION STATEMENT A:

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED

DISTRIBUTION STATEMENT B:

DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES ONLY;
(Indicate Reason and Date Below). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED
TO (Indicate Controlling DoD Office Below).

DISTRIBUTION STATEMENT C:

DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND THEIR CONTRACTORS;
(Indicate Reason and Date Below). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED
TO (Indicate Controlling DoD Office Below).

DISTRIBUTION STATEMENT D:

DISTRIBUTION AUTHORIZED TO DOD AND U.S. DOD CONTRACTORS ONLY; (Indicate Reason
and Date Below). OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office Below).

DISTRIBUTION STATEMENT E:

DISTRIBUTION AUTHORIZED TO DOD COMPONENTS ONLY; (Indicate Reason and Date Below).
OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office Below).

DISTRIBUTION STATEMENT F:

FURTHER DISSEMINATION ONLY AS DIRECTED BY (Indicate Controlling DoD Office and Date
Below) or HIGHER DOD AUTHORITY.

DISTRIBUTION STATEMENT X:

DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND PRIVATE INDIVIDUALS
OR ENTERPRISES ELIGIBLE TO OBTAIN EXPORT-CONTROLLED TECHNICAL DATA IN ACCORDANCE
WITH DOD DIRECTIVE 5230.25, WITHHOLDING OF UNCLASSIFIED TECHNICAL DATA FROM PUBLIC
DISCLOSURE, 6 Nov 1984 (Indicate date of determination). CONTROLLING DOD OFFICE IS (Indicate
Controlling DoD Office).

The cited documents has been reviewed by competent authority and the following distribution statement is
hereby authorized.

A

(Statement)

OFFICE OF NAVAL RESEARCH
CORPORATE PROGRAMS DIVISION
ONR 353
800 NORTH QUINCY STREET
ARLINGTON, VA 22217-5660

(Controlling DoD Office Name)

(Reason)

DEBRA T. HUGHES
DEPUTY DIRECTOR
CORPORATE PROGRAMS OFFICE

(Signature & Typed Name)

(Assigning Office)

(Controlling DoD Office Address,
City, State, Zip)

19 SEP 1995

(Date Statement Assigned)